

Prescription for Cosmic Information Extraction from Multiple Sky Maps

Steven Gratton*

Institute of Astronomy, Madingley Road, Cambridge, CB3 0HA, UK

(Dated: May 1, 2008)

This paper presents a prescription for distilling the information contained in the cosmic microwave background radiation from multiple sky maps that also contain both instrument noise and foreground contaminants. The prescription is well-suited for cosmological parameter estimation and accounts for uncertainties in the cosmic microwave background extraction scheme. The technique is computationally viable at low resolution and may be considered a natural and significant generalization of the “Internal Linear Combination” approach to foreground removal. An important potential application is the analysis of the multi-frequency temperature and polarization data from the forthcoming Planck satellite.

I. INTRODUCTION

The detection and subsequent investigation of the cosmic microwave background (CMB) radiation over the past four decades has been essential to our current understanding of the universe. Initially providing evidence for an early hot dense radiation-dominated phase to the universe [1, 2], the CMB’s spatial distribution is now studied in precise detail (see [3, 4, 5, 6, 7, 8, 9] for the latest analysis of the WMAP satellite data by the WMAP science team) for clues both about the state of the universe at the beginning of the radiation era and the universe’s composition, structure and subsequent evolution. As cosmologists we have been exceedingly fortunate that, from our vantage point of the earth, the galaxy (away from its plane) is both sufficiently transparent and sufficiently lacking in emission for us to be able to readily measure the intensity of the cosmic signal shining through. The cosmic signal is also moderately polarized and this carries important additional cosmological information, in particular about the presence of gravity waves in the universe (which could be generated by say a high energy inflationary phase before the radiation era) and about the reionization of the universe. Unfortunately, with the signal being weaker and the polarization of the galactic emission being less understood than its intensity, both the detection and analysis of the polarization part of the signal are much more challenging. In addition, there is also extragalactic contamination, such as that from point sources. This paper presents a technique to extract the cosmological information out of multiple sky maps, including polarization ones, in the presence of both instrument noise and foregrounds. The technique rests upon the assumption that the foregrounds do not have the same frequency response as the CMB and requires that the sky maps probe linearly independent parts of the spectrum of the sky signal.

The prescription presented here relates to and builds upon a number of approaches already in the literature. The COBE team considered three approaches to separating galactic emission from the cosmic signal, one involving modelling known emissions using non-CMB data, another involving fitting the maps to functions of given spectral index, and a final one involving linearly combining their multifrequency maps to cancel the dominant galactic emission [10]. The WMAP team primarily use a template subtraction procedure to mitigate the effects of foreground contaminants in power spectrum estimation, and study the foregrounds themselves via maximum entropy and most recently Markov Chain Monte Carlo parametric methods. However, initially for visualization purposes but later also for analysis they additionally introduced the “Internal Linear Combination” (ILC) scheme, forming a linear combination of their sky maps and choosing the weights to minimize the variance between the maps whilst being constrained to preserve unit response to the CMB [4, 11, 12]. An harmonic mode-by-mode equivalent of ILC was presented in [13] and applied to the WMAP data in [14]. An alternative harmonic-based generalization of the ILC technique was recently presented in [15]. The “Independent Component Analysis” (ICA) signal processing technique (see [16]), which attempts to use non-gaussianity of the one-pixel distribution to separate data into independent signals, has been tested for cosmological uses on the COBE, BEAST and WMAP data [17, 18, 19, 20]. A possible weakness of ICA to CMB extraction is that the CMB is believed to be very close to gaussian and so can only emerge as what is left behind after the non-gaussian foregrounds have been removed. The related “Correlated Component Analysis” idea uses pixel-pixel cross correlations instead of non-quadratic single-pixel statistics to separate signals and has also been applied to the WMAP data [21, 22]. A modification of ICA that forces it to take into account the black body nature of the CMB (a key feature of the approach described here) was recently presented in [23]. The “spectral matching” approach

*Electronic address: stg20@cam.ac.uk

of [24, 25] shares many similarities to that presented here, and a recent paper [26] presents an “additive component”-based separation technique. [27] suggests fitting model parameters at low resolution and then using these to solve for high resolution maps; see also the very recent work [28] for more on parametric component separation. A Gibbs sampling based approach to component separation and CMB power spectrum estimation was presented in [29] and applied to the WMAP data in [30]. The WMAP team also test this approach in [6].

Unlike many of the above papers, this work focusses on likelihood estimation for cosmological models as opposed to CMB sky map production. This requires a quantification of the uncertainties related to the CMB extraction (as also stressed by Ref. [31]). Of course this is a somewhat ill-defined problem seeing as one does not know precisely what the foregrounds are. By putting relatively weak priors on the foregrounds though one might hope that such errors are being estimated if anything conservatively. Our prescription allows one to naturally incorporate non-CMB datasets and the information they contain about the foregrounds into the analysis. With the imminent launch of the Planck satellite [32] and the significant new information on the polarization of the CMB that it should deliver, this work is also notable in treating all Stokes parameters of the CMB in a unified manner. Numerical testing of the scheme and application to existing WMAP data are underway.

II. DATA AND PRIORS

Our starting point shall be a collection of n sky maps. Any that are usefully described in terms of a temperature will be assumed to be in thermodynamic temperature units. Note from the start that these maps do not all have to be “CMB” maps; other data sets (e.g. radio surveys, starlight polarization maps, point source maps, the WMAP “spurious signal” maps) can be included in the analysis in a unified manner and might be useful if they have physical correlations with contaminants in the CMB channels.

We assume the sky map is discretized into elements, typically pixels or spherical harmonic coefficients up to some l_{\max} , but perhaps say wavelet coefficients. For each element i we have the associated measurements for each sky map. We can stack these measurements into a vector $\mathbf{X}(i)$. These vectors can be further stacked into a big vector \mathbf{X} (the entire data set). We’ll assume we can estimate or calculate the inverse noise covariance matrix \mathbf{N}^{-1} for \mathbf{X} .

Next, let us assume a linear relation between \mathbf{X} and some underlying “signals” \mathbf{S} , i.e.

$$\mathbf{X} = \mathbf{A}\mathbf{S} + \mathbf{M}, \quad (1)$$

where \mathbf{M} is the noise realization. Some of the signals will of course be the CMB, and the others will be the foregrounds. These “effective” foreground signals don’t necessarily have to be thought of as physical processes, just as unwanted contaminants. We shall not assume that these foregrounds are uncorrelated with each other, but they will, however, almost by definition be taken to independent of the CMB. To extract cosmological information, we shall work towards a probability distribution for that part of \mathbf{S} associated with the CMB sky.

A key assumption shall be that the CMB is a blackbody. As a consequence, the “foreground” signals will implicitly be assumed to be linearly independent of the CMB in frequency space.

We’ll typically assume that the “mixing matrix” \mathbf{A} is block-diagonal in pixel space, meaning that measurements in a given direction only depend on the signals in that direction (and we are assuming that any beam convolution effects have already been accounted for). In this case, we can write:

$$\mathbf{X}(i) = \mathbf{A}(ii)\mathbf{S}(i) + \mathbf{M}(i). \quad (2)$$

(Here and onwards, a matrix followed by element indices denote the submatrix appropriate to the indices of the original matrix.) $\mathbf{S}(i)$ encodes the strength of the signals, and $\mathbf{A}(ii)$ controls how much each signal affects each instrument channel. The degeneracy between the amplitude of a column of $\mathbf{A}(ii)$ and the normalization of a given foreground signal is in principle handled by assumed priors, and the discrete symmetry of column interchange of $\mathbf{A}(ii)$ leads to an uniform multiplicative overcounting that can be ignored.

The blackbody nature of the CMB is realized as a requirement on the form of the $\mathbf{A}(ii)$ ’s. If $\mathbf{S}(i)^T$ takes the form $(\mathbf{S}_{\text{CMB}}(i)^T, \mathbf{F}(i)^T)$, with \mathbf{S}_{CMB} for the CMB component(s) and \mathbf{F} for the foregrounds, then $\mathbf{A}(ii)$ must take the form

$$\mathbf{A}(ii) = \left(\begin{array}{c|cc} \mathbf{e} & ? & ? \\ \hline \vdots & ? & ? \end{array} \right), \quad (3)$$

where \mathbf{e} is a “tall” matrix whose width equals the number of components of $\mathbf{S}_{\text{CMB}}(i)$ and whose elements are ones or zeroes (since temperature-type maps are assumed to be in thermodynamic temperature units). If a channel is polarization-sensitive, then its corresponding part of \mathbf{e} will be a “small” identity matrix (with \mathbf{i} denoting identity matrices of various sizes from now on). If a channel is independent of the CMB its corresponding part of \mathbf{e} will be a

“small” (row) matrix of zeroes (all “zero” matrices will be denoted $\mathbf{0}$). Such a channel is still useful in the analysis because it potentially responds to some of the foreground signals and thus carries information about them that can be used in extracting the CMB. The part of \mathbf{a} marked ? encodes the spectral indices of the foreground signals.

Because some physical foregrounds (e.g. galactic synchrotron radiation) are known to have a spectral index that varies across the sky, one might expect that our $\mathbf{A}(ii)$ should indeed be a function of the element i . However, an alternative procedure is possible if we have many frequency channels; we can take $\mathbf{A}(ii) = \mathbf{a}$ for all i , and describe a single physical foreground with varying spectral index as two or more correlated effective signal components which have linearly independent corresponding columns of \mathbf{a} .

From now on we will take \mathbf{a} to be a square n -by- n matrix, assuming there to be as many signal components as we have frequency channels. As we’ll be marginalizing out the foreground signals, it seems reasonable to allow for as many of them as we might need to describe the data on top of the noise.

We are aiming towards using Bayes’ Theorem to get a probability for a CMB model (or perhaps a CMB sky map) given the data, marginalized over foregrounds. Let us set down the relevant conditional probabilities that we shall need. Firstly,

$$p(\mathbf{S}|\text{CMB model, f.g. model}) = p(\mathbf{S}_{\text{CMB}}|\text{CMB model}).p(\mathbf{F}|\text{f.g. model}) \quad (4)$$

by the independence assumption, where “f.g.” is short for foreground. Here \mathbf{S}_{CMB} is the CMB signal (temperature and perhaps polarization) and \mathbf{F} are the foreground signals, as introduced above. Typically the CMB model is describable in terms of the CMB power spectra, but in principle we could consider non-gaussian corrections too. In the gaussian case,

$$p(\mathbf{S}_{\text{CMB}}|\text{CMB model}) = \frac{1}{\sqrt{|2\pi\mathbf{C}|}} e^{-\mathbf{S}_{\text{CMB}}^T \mathbf{C}^{-1} \mathbf{S}_{\text{CMB}}/2}. \quad (5)$$

For isotropic models in harmonic space, \mathbf{C} is diagonal in l and conventionally expressed in terms of C_l ’s (i.e. $\langle a_{lm}^P a_{l'm'}^Q \rangle = C_l^{PQ} \delta_{l,l'} \delta_{m,-m'}$ where P and Q denote Stokes parameter type.).

Next, the probability of the data given the signals and the mixing matrix is given by the probability of obtaining the noise realization $\mathbf{M} = \mathbf{X} - \mathbf{A}\mathbf{S}$:

$$p(\mathbf{X}|\mathbf{S}, \mathbf{A}) = \frac{1}{\sqrt{|2\pi\mathbf{N}|}} e^{-(\mathbf{X}-\mathbf{A}\mathbf{S})^T \mathbf{N}^{-1} (\mathbf{X}-\mathbf{A}\mathbf{S})/2}. \quad (6)$$

Here we have used maximum entropy to assign a gaussian distribution to the noise (see [33]) and have decided not to add in and marginalize over any non-gaussian corrections we could imagine (which would complicate the analysis). In some cases we might wish to consider “inverse noise matrices” \mathbf{N}^{-1} that are not actually the inverse of anything due to having had say monopole and dipole components, point sources, regions of strong galactic emission or instrumental effects projected out of them. (Such projections ensure that a model is not penalized if the data and model “disagree” along such directions.) In this case the determinant factor will not exist but seeing as it is independent of the model it can be safely ignored. Note that in many applications where the sky map data is actually derived with a maximum-likelihood approach from a timestream it is in fact the inverse noise covariance matrix and not the noise covariance matrix itself that naturally emerges.

Now we need some priors. For the C_l ’s, we might employ a Jeffreys’-type prior l by l , or alternatively assume they come from a known cosmological model which itself has priors on its parameters. In a middle-ground position we might incorporate a positive correlation between neighbouring C_l ’s, expecting them to be given by integrals over related transfer functions times the same underlying three-dimensional primordial power spectrum [34]. As simple test cases we could assume that the CMB is gaussian white noise on the sky with some unknown amplitude or that the CMB is scale-invariant with unknown amplitude. We shall see that the “white noise” assumption leads directly to the ILC approach as used by the WMAP team.

We’ve already effectively imposed a prior that the mixing matrix \mathbf{A} is block-diagonal in pixel space and have argued that it can be taken to be isotropic across the sky if we have enough sky maps to work with. (If we do relax the isotropy assumption on \mathbf{a} then we should parametrize a slow variation of \mathbf{a} across the sky with a suitable probability distribution over the parameters; see the recent work of [15] for an application of this type of approach to ILC map making.) One might be concerned whether it is most natural to put a measure on \mathbf{a} or on its inverse, since the former is perhaps more natural when thinking about the spectral response of physical foregrounds whereas the latter is heuristically at least what is needed to go from the data to the CMB. In fact we shall develop a measure that is invariant under inversion. A matrix measure naturally involves some power of the determinant of the matrix,

and from $\mathbf{a}^{-1}\mathbf{a} = \mathbf{i}$ one can see that $\delta\mathbf{a}^{-1} = -\mathbf{a}^{-1}\delta\mathbf{a}\mathbf{a}^{-1}$ and hence that $\partial(\mathbf{a}^{-1})/\partial(\mathbf{a}) = |\mathbf{a}^{-1}|^{2n}$.¹ So a measure of $|\mathbf{a}|^m d\mathbf{a}$ transforms into a measure $|\mathbf{a}^{-1}|^{-m-2n} d\mathbf{a}^{-1}$ for some power m . Full form symmetry is achieved if we choose $m = -m - 2n$, or $m = -n$. This choice yields the Haar measure on $GL(n, R)$, the group of symmetries of a real n -dimensional vector space, and so seems very natural. This measure is also scale-invariant as well as inversion-invariant. From this base we now consider our constraints on the mixing matrix from our taking the CMB to be black-body and to be independent from the foregrounds. Start from the matrix equation $\mathbf{a}^{-1}\mathbf{a} = \mathbf{i}$. If we write \mathbf{a}^{-1} as

$$\mathbf{a}^{-1} = \begin{pmatrix} \mathbf{w}^T & \dots \\ \mathbf{u}^T & \dots \end{pmatrix}, \quad (7)$$

then we must have

$$\mathbf{w}^T \mathbf{e} = \mathbf{i}, \quad (8a)$$

$$\mathbf{u}^T \mathbf{e} = \mathbf{0}, \quad (8b)$$

where \mathbf{w} and \mathbf{u} are appropriate “tall” matrices. From this we can see that delta functions required in our measure on \mathbf{a} to force the first columns of \mathbf{a} to equal \mathbf{e} correspond to the term $|\mathbf{a}^{-1}|^{-1} \delta(\mathbf{w}^T \mathbf{e} - \mathbf{i}) \delta(\mathbf{u}^T \mathbf{e})$ in a measure for \mathbf{a}^{-1} . With the form of \mathbf{e} as discussed above, Eqs. (8) render $|\mathbf{a}^{-1}|$ independent of \mathbf{w} .² Hence the components \mathbf{w} needed to extract the CMB only actually appear in the delta function $\delta(\mathbf{w}^T \mathbf{e} - \mathbf{i})$ and not in any prefactor in a base measure on \mathbf{a}^{-1} .³

The final prior is for the foregrounds. We shall take them to be gaussian, with auto- and cross- l by l correlations. The gaussian assumption is not ideal; indeed, the WMAP K- and KQ- sequence of masks are actually constructed by excluding pixels whose temperatures fall into the skewed high tail of the one-point temperature distribution function. However, by taking the inverse covariance to zero, a gaussian model does at least provide an analytically-controlled method of approaching a flat prior on the foregrounds. A flat prior might be particularly appropriate for polarization, given our limited physical understanding of galactic polarization, and so is what a maximum entropy argument would yield.⁴

III. PROBABILITIES FOR CMB POWER SPECTRA WITH NOISE

We now move on to our main interest, the calculation of probabilities for cosmological models, expressed in terms of their predicted CMB power spectra, in presence of general (anisotropic and correlated) instrument noise and foreground contaminants.

We shall take a limit towards a flat prior on the foregrounds, starting from a gaussian prior on them. As we shall see, this limit conveniently decouples the foreground-related parts of the inverse mixing matrix from the rest of the problem. We’ll call the “grand” covariance matrix, of CMB and foregrounds combined, \mathbf{G} , and under our stated assumptions this fully specifies our CMB and foreground models.

We have:

$$p(\mathbf{X}|\mathbf{a}^{-1}, \mathbf{G}) = \int d\mathbf{S} p(\mathbf{X}, \mathbf{S}|\mathbf{a}^{-1}, \mathbf{G}) \quad (9)$$

$$= \int d\mathbf{S} p(\mathbf{X}|\mathbf{S}, \mathbf{a}^{-1}, \mathbf{G}) p(\mathbf{S}|\mathbf{a}^{-1}, \mathbf{G}) \quad (10)$$

$$\propto \int d\mathbf{S} \frac{e^{-(\mathbf{X}-\mathbf{AS})^T \mathbf{N}^{-1} (\mathbf{X}-\mathbf{AS})/2}}{\sqrt{|2\pi\mathbf{N}|}} \sqrt{|\mathbf{G}^{-1}/2\pi|} e^{-\mathbf{S}^T \mathbf{G}^{-1} \mathbf{S}/2} \quad (11)$$

¹ The latter is most easily seen by considering the two linear transformations on $\delta\mathbf{a}$ that aggregate to give $\delta\mathbf{a}^{-1}$, i.e. bracketing the latter as $-(\mathbf{a}^{-1}\delta\mathbf{a}\mathbf{a}^{-1})\mathbf{a}^{-1}$, using the chain rule for Jacobians and performing row and column interchange operations on each resulting n^2 -by- n^2 matrix.

² This can be seen by using Eqs. (8) to substitute for the variables appearing in the first n_{CMB} columns of \mathbf{a}^{-1} , then appropriately adding in the other columns to make the first ones equal $(\mathbf{i}, \mathbf{0})^T$ in the determinant calculation.

³ If we do have useful information on the foreground spectra we would presumably first encode this in \mathbf{a} and then consider transforming to \mathbf{a}^{-1} . If this information is actually coming from other non-CMB sky maps, it might perhaps be easier to perform a combined analysis including these maps as discussed above but with a larger \mathbf{a} .

⁴ As with for the mixing matrix, we could take any data sets that do give us extra information into account by directly including them in the analysis.

and to do the integral over \mathbf{S} it is useful to change variables to $\mathbf{U} \equiv \mathbf{A}\mathbf{S}$, generating a $|\mathbf{A}|$ determinant factor. With the mixing matrix being uniform across the sky, \mathbf{A} decomposes into \mathbf{a} 's down the diagonal and so $|\mathbf{A}| = |\mathbf{a}|^{N_{\text{elem}}}$. Note that the right-hand-side has been written so as to not explicitly depend on \mathbf{G} , only on \mathbf{G}^{-1} , to make the limiting process clearer. We can perform the gaussian integral over \mathbf{S} and use the result in Bayes' theorem, in conjunction with our priors discussed in Sec. II, to find:

$$p(\mathbf{a}^{-1}, \mathbf{G}|\mathbf{X}) \propto p(\mathbf{X}|\mathbf{a}^{-1}, \mathbf{G})p(\mathbf{a}^{-1}, \mathbf{G}) \quad (12)$$

$$\propto p(\mathbf{G}) \frac{\sqrt{|\mathbf{G}^{-1}|}}{|\mathbf{a}^{-1}|^{-N_{\text{elem}}}} \frac{\delta(\mathbf{w}^T \mathbf{e} - \mathbf{i})\delta(\mathbf{u}^T \mathbf{e})}{|\mathbf{a}^{-1}|^{n+1}} \frac{e^{+\mathbf{X}^T \mathbf{N}^{-1}(\mathbf{N}^{-1} + \mathbf{A}^{-T} \mathbf{G}^{-1} \mathbf{A}^{-1})^{-1} \mathbf{N}^{-1} \mathbf{X}/2 - \mathbf{X}^T \mathbf{N}^{-1} \mathbf{X}/2}}{\sqrt{|\mathbf{N}^{-1} + \mathbf{A}^{-T} \mathbf{G}^{-1} \mathbf{A}^{-1}|}}. \quad (13)$$

To do this integral we needed $\mathbf{N}^{-1} + \mathbf{A}^{-T} \mathbf{G}^{-1} \mathbf{A}^{-1}$ to be invertible. We also see that as long as the number of elements is much larger than the number of frequency channels, the $|\mathbf{a}^{-1}|^{n+1}$ term from our prior on \mathbf{a}^{-1} becomes insignificant and could be ignored relative to the $|\mathbf{a}^{-1}|^{N_{\text{elem}}}$ term from the likelihood. We keep both however, defining $N \equiv N_{\text{elem}} - n - 1$.

It is convenient to take the flat-foreground-prior limit $\mathbf{F}^{-1} \rightarrow 0$ at this point. First note that $|\mathbf{G}^{-1}| = |\mathbf{C}^{-1}||\mathbf{F}^{-1}|$ since CMB and foregrounds are assumed independent. By taking the limit in the same manner for the different CMB models that we are considering, we can safely ignore the vanishing $|\mathbf{F}^{-1}|$ factor (alternatively we can say that we are only interested in likelihood ratios between models, in which case the factor vanishes).

Next, we consider $\mathbf{A}^{-T} \mathbf{G}^{-1} \mathbf{A}^{-1}$. Written out more explicitly, this takes the form

$$\begin{pmatrix} \mathbf{a}^{-T} \mathbf{G}^{-1}(1,1) \mathbf{a}^{-1} & \mathbf{a}^{-T} \mathbf{G}^{-1}(1,2) \mathbf{a}^{-1} & & \\ \mathbf{a}^{-T} \mathbf{G}^{-1}(2,1) \mathbf{a}^{-1} & \mathbf{a}^{-T} \mathbf{G}^{-1}(2,2) \mathbf{a}^{-1} & & \\ & & \ddots & \end{pmatrix} \quad (14)$$

where we have decomposed \mathbf{G}^{-1} into element-labelled blocks, $\mathbf{G}^{-1}(ij)$ corresponding to the block of \mathbf{G}^{-1} relating to elements i and j . As $\mathbf{F}^{-1} \rightarrow 0$, $\mathbf{a}^{-T} \mathbf{G}^{-1}(ij) \mathbf{a}^{-1} \rightarrow \mathbf{w} \mathbf{C}^{-1}(ij) \mathbf{w}^T$. So, introducing the nN_{elem} -by- nN_{elem} matrix \mathbf{Q} :

$$\mathbf{Q} \equiv \begin{pmatrix} \mathbf{w} & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{w} & \mathbf{0} & \\ \vdots & \vdots & \vdots & \end{pmatrix}, \quad (15)$$

we have

$$\mathbf{A}^{-T} \mathbf{G}^{-1} \mathbf{A}^{-1} \rightarrow \mathbf{Q} \mathbf{C}^{-1} \mathbf{Q}^T. \quad (16)$$

Thus the probability factors into a term depending on \mathbf{w} and \mathbf{C} alone and a term depending on \mathbf{u} alone (the $\delta(\mathbf{u}^T \mathbf{e})/|\mathbf{a}^{-1}|^{-N}$ part). When we integrate over \mathbf{a} , or equivalently over \mathbf{w} and \mathbf{u} , to obtain the marginalized probability for the CMB model, the \mathbf{u} integral thus factors off to give a multiplicative constant that can be ignored.

Hence we obtain an effective probability for a CMB model of

$$p(\mathbf{C}|\mathbf{X}) \propto \int d\mathbf{w} \delta(\mathbf{w}^T \mathbf{e} - \mathbf{i}) \frac{e^{+\mathbf{X}^T \mathbf{N}^{-1}(\mathbf{N}^{-1} + \mathbf{Q} \mathbf{C}^{-1} \mathbf{Q}^T)^{-1} \mathbf{N}^{-1} \mathbf{X}/2 - \mathbf{X}^T \mathbf{N}^{-1} \mathbf{X}/2}}{\sqrt{|\mathbf{C}||\mathbf{N}^{-1} + \mathbf{Q} \mathbf{C}^{-1} \mathbf{Q}^T|}}. \quad (17)$$

The integral over \mathbf{w} will be performed using the saddle-point method developed in Appendix A.

Before integrating though, there may exist opportunities to significantly simplify (17), depending on the invertibility properties of the matrices remaining. If so, the size of subsequent matrix operations is substantially reduced, easing the computational burden. For the rest of this section we consider the simplest case of both \mathbf{N}^{-1} and \mathbf{C}^{-1} being invertible; more complicated cases are deferred to subsections below. Following the discussion in Appendix C, we are able to apply formulae (C4) and (C3) to simplify (17) to

$$p(\mathbf{C}|\mathbf{X}) \propto p(\mathbf{C}) \int d\mathbf{w} \delta(\mathbf{w}^T \mathbf{e} - \mathbf{i}) \frac{e^{-\mathbf{X}^T \mathbf{Q}(\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q})^{-1} \mathbf{Q}^T \mathbf{X}/2}}{\sqrt{|\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q}|}}. \quad (18)$$

Note how the size of the matrix that needs to be inverted and have its determinant found has reduced from nN_{elem} -by- nN_{elem} to $n_{\text{CMB}}N_{\text{elem}}$ -by- $n_{\text{CMB}}N_{\text{elem}}$, where n_{CMB} is the number of CMB Stokes parameters under consideration, saving $O((n/n_{\text{CMB}})^3)$ in required computation.

Comparing with Eq. (A1), we have

$$S = \frac{1}{2} (\ln |\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q}| + \mathbf{X}^T \mathbf{Q} (\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q})^{-1} \mathbf{Q}^T \mathbf{X}), \quad (19)$$

and by varying with respect to \mathbf{w} the appropriate first and second derivatives for S can be read off for use in (A11):

$$\delta_{\mathbf{w}} S = \text{tr} \delta \mathbf{w}^T \{ \mathbf{N}(ij) \mathbf{w} (\mathbf{D}^{-1}(ji) - \mathbf{D}^{-1}(jm) \mathbf{X}(m) \mathbf{X}^T(n) \mathbf{D}^{-1}(ni)) + \mathbf{X}(i) \mathbf{X}^T(j) \mathbf{w} \mathbf{D}^{-1}(ji) \} \quad (20)$$

$$\begin{aligned} \delta_{\mathbf{w}}^2 S = & \text{tr} \delta \mathbf{w}^T \mathbf{N}(ij) \delta \mathbf{w} (\mathbf{D}^{-1}(ji) - \mathbf{D}^{-1}(jm) \mathbf{X}(m) \mathbf{X}^T(n) \mathbf{D}^{-1}(ni)) / 2 \\ & + \text{tr} \delta \mathbf{w}^T \mathbf{N}(ij) \mathbf{w} \mathbf{D}^{-1}(jk) \delta \mathbf{w}^T \mathbf{N}(ki) \mathbf{w} \\ & + \text{tr} \delta \mathbf{w}^T \mathbf{N}(ij) \mathbf{w} \mathbf{D}^{-1}(jk) \mathbf{w}^T \mathbf{N}(ki) \delta \mathbf{w} \\ & + \text{tr} \delta \mathbf{w}^T \mathbf{N}(ij) \mathbf{X}(i) \mathbf{X}^T(j) \mathbf{w} \mathbf{D}^{-1}(ji) / 2 \\ & - \text{tr} \delta \mathbf{w}^T \mathbf{X}(i) \mathbf{X}^T(j) \mathbf{w} \mathbf{D}^{-1}(jk) (\delta \mathbf{w}^T \mathbf{N}(km) \mathbf{w} + \mathbf{w}^T \mathbf{N}(mn) \delta \mathbf{w}) \mathbf{D}^{-1}(ni) \\ & + \text{tr} \delta \mathbf{w}^T \mathbf{N}(ij) \mathbf{w} \mathbf{D}^{-1}(jk) \delta \mathbf{w}^T \mathbf{N}(km) \mathbf{w} \mathbf{D}^{-1}(mn) \mathbf{w}^T \mathbf{X}(n) \mathbf{X}^T(p) \mathbf{w} \mathbf{D}^{-1}(pi) \\ & + \text{tr} \delta \mathbf{w}^T \mathbf{N}(ij) \mathbf{w} \mathbf{D}^{-1}(jk) \mathbf{w}^T \mathbf{N}(km) \delta \mathbf{w} \mathbf{D}^{-1}(mn) \mathbf{w}^T \mathbf{X}(n) \mathbf{X}^T(p) \mathbf{w} \mathbf{D}^{-1}(pi) / 2 \\ & + \text{tr} \delta \mathbf{w}^T \mathbf{D}^{-1}(ij) \delta \mathbf{w} \mathbf{N}(jk) \mathbf{w} \mathbf{D}^{-1}(km) \mathbf{w}^T \mathbf{X}(m) \mathbf{X}^T(n) \mathbf{w} \mathbf{D}^{-1}(np) \mathbf{w}^T \mathbf{N}(pi) / 2 \end{aligned} \quad (21)$$

where we have defined \mathbf{D} to equal $\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q}$ and assumed implied summation over repeated element indices. (To obtain full equivalence with Eq. (A11) we need to map from the matrix \mathbf{w} here to the vector x there. This can be done explicitly by “vectorization” of \mathbf{w} if required, as discussed in Appendix B.) The numerical calculation time is dominated by the required $O((n_{\text{CMB}} N_{\text{elem}})^3)$ inversion of \mathbf{D} . So as long as the initial guess for \mathbf{w} is reasonable⁵, one should be able to converge to the saddle-point expression for \mathbf{w} in a few $O((n_{\text{CMB}} N_{\text{elem}})^3)$ steps. Having converged to a saddle point solution $\mathbf{w}_{\text{s.p.}}$, we evaluate the exponent (19) there, calculate a prefactor⁶ involving second derivatives of S at the saddle point as discussed in Appendix A and multiply by any prior in order to finally obtain the posterior probability for the CMB model in question.

A. Noise Matrices with Projections

Often one wishes to project additional degrees of freedom out of a formally invertible inverse noise matrix in order to render our posterior probabilities insensitive to certain complications with the data that would otherwise be unaccounted for. For example, one might project out monopole and dipole contributions of the maps from an experiment like WMAP, the monopole because the experiment is basically differential and the dipole because of the earth’s motion relative to the CMB. Or one might project out regions of strong emission from the galaxy because one does not expect the relatively simple foreground model used to be accurate there. The WMAP team also project out a “transmission loss imbalance” mode from their maps. In some cases one can sometimes just reduce the dimensions of the matrices involved; e.g. one might just use the $l > 1$ modes in harmonic space to forget about the monopole and dipole, or for a galactic cut working in pixel space one might simply only consider data from pixels outside of the cut. However, when say both a monopole and dipole projection and a galactic cut are made, no such simple approach is possible, and one must proceed as follows.

Each of the n_{proj} modes to be projected out are expressed as a column vector, and these column vectors are arranged into an $n N_{\text{elem}}$ -by- n_{proj} matrix \mathbf{P} . Then, \mathbf{N}^{-1} is replaced by:

$$\mathbf{N}^{-1} \rightarrow \overline{\mathbf{N}^{-1}}|_{\mathbf{P}} \equiv \mathbf{N}^{-1} - \mathbf{N}^{-1} \mathbf{P} (\mathbf{P}^T \mathbf{N}^{-1} \mathbf{P})^{-1} \mathbf{P}^T \mathbf{N}^{-1}. \quad (22)$$

(Note that the normalization of the modes cancels out in the formation of $\overline{\mathbf{N}^{-1}}|_{\mathbf{P}}$.) $\overline{\mathbf{N}^{-1}}|_{\mathbf{P}}$ has been constructed so that when it is multiplied into any linear combination of the modes to be projected out it returns zero. Hence any discrepancy between the data and a potential signal along these modes is not penalized in the likelihood.

⁵ A suitable starting place should be obtainable from the “noise-free” solution evaluated for a reasonable fiducial model, as derived in Sec. IV.

⁶ This prefactor coming from the integral over \mathbf{w} is helping to encode the uncertainties from the foreground separation into the likelihood; this term would be missed in any approach that attempts to use a standard CMB likelihood formula applied to some best-fit linearly combined map and effective noise matrix.

Eq. (17) becomes

$$p(\mathbf{C}|\mathbf{X}) \propto \int d\mathbf{w} \delta(\mathbf{w}^T \mathbf{e} - \mathbf{i}) \frac{e^{+\mathbf{X}^T \overline{\mathbf{N}^{-1}}|_{\mathbf{P}} (\overline{\mathbf{N}^{-1}}|_{\mathbf{P}} + \mathbf{Q} \mathbf{C}^{-1} \mathbf{Q}^T)^{-1} \overline{\mathbf{N}^{-1}}|_{\mathbf{P}} \mathbf{X}/2 - \mathbf{X}^T \overline{\mathbf{N}^{-1}}|_{\mathbf{P}} \mathbf{X}/2}}{\sqrt{|\mathbf{C}| |(\overline{\mathbf{N}^{-1}}|_{\mathbf{P}} + \mathbf{Q} \mathbf{C}^{-1} \mathbf{Q}^T)|}}, \quad (23)$$

which, after some work involving (C4) and (C3) and keeping \mathbf{N} to refer to the assumed-invertible formal noise matrix, reduces to:

$$p(\mathbf{C}|\mathbf{X}) \propto p(\mathbf{C}) \int d\mathbf{w} \delta(\mathbf{w}^T \mathbf{e} - \mathbf{i}) \frac{e^{-\mathbf{X}^T \mathbf{Q} (\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q})^{-1} |_{\mathbf{Q}^T \mathbf{P}} \mathbf{Q}^T \mathbf{X}/2}}{\sqrt{|\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q}| |\mathbf{P}^T \mathbf{Q} (\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q})^{-1} \mathbf{Q}^T \mathbf{P}|}}. \quad (24)$$

We can now expand the delta function as a gaussian and integrate as above. Additional terms in the first and second derivatives of the exponent come from the projection, further complicating the formulae, but the numerical time of the calculation is still dominated by the inversion of $\mathbf{C} + \mathbf{Q}^T \mathbf{N} \mathbf{Q}$.

There is some subtlety about which projections can be handled in this manner. For example, with temperature maps alone, one cannot independently project out all measurements associated with a given element; the second determinant under the square root turns out to be proportional to $|\mathbf{w} \mathbf{w}^T|^{N_{\text{elem}}}$ and is thus singular. (This is happening perhaps because as the foregrounds are being marginalized out and the dimensions effectively reducing from $n N_{\text{elem}}$ to N_{elem} the previously orthogonal projections might be becoming linearly dependent.) Potentially including polarisation again, one can however project out any response to the CMB in the i 'th element say with $\mathbf{P}(i)$ equal to \mathbf{e} and all other entries in \mathbf{P} zero. With the help of the delta function $\mathbf{Q}^T \mathbf{P}(i)$ becomes equal to the identity matrix \mathbf{i} . Finally, one obtains just the result one would have gotten by forgetting about the element in question in the first place and starting with a reduced problem with only $N_{\text{elem}} - 1$ elements rather than N_{elem} elements, as can be seen using rules for the determinants of blocked matrices and their inverses.⁷ By combining many such projections into a wide projection matrix one can form a mask, or by considering an appropriately weighted vector sum of such projections one can project out spatially coherent modes such as the monopole or dipole from pixellized data.

B. The general case

Finally we mention the case when the inverse noise matrix is completely general. Now, we cannot use the Woodbury formula to simplify the exponential in (17). Rather, we have to use the effective exponent from (17) as is and directly proceed with the replacement of the delta function with a gaussian. It is in fact very appealing to work like this directly with inverse noise matrices and inverse-noise-weighted maps, as these are what naturally emerge out of a maximum-likelihood analysis of timestream data. However, the practical cost is that one now has to invert full $n N_{\text{elem}}$ -by- $n N_{\text{elem}}$ matrices for each CMB model considered.

IV. THE NO-NOISE LIMIT

It is interesting to consider what happens if the noise is negligible. Starting from (18), we can take the limit $\mathbf{N} \rightarrow 0$ to obtain:

$$p(\mathbf{C}|\mathbf{X})_{\text{no noise}} \propto \int d\mathbf{w} \delta(\mathbf{w}^T \mathbf{e} - \mathbf{i}) \frac{p(\mathbf{C})}{\sqrt{|\mathbf{C}|}} e^{-S'}, \quad (25)$$

with exponent $S' = \text{vec}(\mathbf{w})^T \hat{\mathbf{R}} \text{vec}(\mathbf{w})/2$ (with “vec” indicating vectorization as discussed in Appendix B, and using Eq. (B5)) where we have defined

$$\hat{\mathbf{R}}(\mathbf{C}) = \sum_{ij} \mathbf{C}^{-1}(ij) \otimes \mathbf{X}(i) \mathbf{X}^T(j). \quad (26)$$

⁷ By only using certain columns of \mathbf{e} in $\mathbf{P}(i)$ one can in fact project out individual Stokes components of the CMB, paving the way for independent temperature and polarization masks.

Here \otimes denotes the Kronecker product, discussed in Appendix B. Notice that S' is now quadratic in the components of \mathbf{w} and hence the saddle point integration over \mathbf{w} is exact. Re-expressing the $\mathbf{w}^T \mathbf{e} = \mathbf{i}$ constraint as

$$(\mathbf{i} \otimes \mathbf{e}^T) \text{vec}(\mathbf{w}) = \text{vec}(\mathbf{i}), \quad (27)$$

Eq. (A12) gives the saddle point for $\text{vec}(\mathbf{w})$ at:

$$\text{vec}(\mathbf{w})_{\text{s.p.}} = \hat{\mathbf{R}}^{-1}(\mathbf{i} \otimes \mathbf{e}) \left((\mathbf{i} \otimes \mathbf{e}^T) \hat{\mathbf{R}}^{-1}(\mathbf{i} \otimes \mathbf{e}) \right)^{-1} \text{vec}(\mathbf{i}) \quad (28)$$

(“s.p.” for saddle point). Using (A13), the posterior for the CMB model is

$$p(\mathbf{C}|\mathbf{X}) \propto \frac{p(\mathbf{C})}{\sqrt{|\mathbf{C}|}} \frac{1}{\sqrt{|\hat{\mathbf{R}}| |(\mathbf{i} \otimes \mathbf{e}^T) \hat{\mathbf{R}}^{-1}(\mathbf{i} \otimes \mathbf{e})|}} e^{-\text{vec}(\mathbf{i})^T ((\mathbf{i} \otimes \mathbf{e}^T) \hat{\mathbf{R}}^{-1}(\mathbf{i} \otimes \mathbf{e}))^{-1} \text{vec}(\mathbf{i})/2}. \quad (29)$$

Eq (28) bears a marked resemblance to the formula for the weights in the ILC approach and this will be expanded upon in the following section.

V. DERIVING SKY MAPS AND RELATIONS WITH THE ILC PROCEDURE

So far in this paper we have concentrated on deriving likelihood functions for CMB models rather than on producing CMB sky maps. However, our procedure can of course be used to generate sky maps and associated noise covariances including contributions from foreground removal.

Now $p(\mathbf{S}|\mathbf{X}) = \int d\mathbf{A} d\mathbf{G} p(\mathbf{S}, \mathbf{A}, \mathbf{G}|\mathbf{X}) \propto \int d\mathbf{A} d\mathbf{G} p(\mathbf{X}|\mathbf{S}, \mathbf{A}, \mathbf{G}) p(\mathbf{S}, \mathbf{A}, \mathbf{G})$ using Bayes' Theorem. Using the same conditional probabilities and priors as above, again introducing $\mathbf{U} = \mathbf{A}\mathbf{S}$, and using Eq. (A8), we obtain:

$$\langle \mathbf{S}_{\text{CMB}}(i) \rangle \approx \int d\mathbf{C} p(\mathbf{C}|\mathbf{X}) \mathbf{w}_{\text{s.p.}}^T(\mathbf{C}) \mathbf{X}(i) \quad (30)$$

along with a somewhat messier expression for $\langle \mathbf{S}_{\text{CMB}}(i) \mathbf{S}_{\text{CMB}}^T(j) \rangle$ using Eq. (A9). In the no-noise limit, these expressions become exact and $\mathbf{w}_{\text{s.p.}}^T(\mathbf{C})$ and $p(\mathbf{C}|\mathbf{X})$ are obtainable from Eqs. (28) and (29) respectively.

To establish a direct link with ILC let us try to construct the CMB temperature (and possibly polarization) maps, \mathbf{S}_{CMB} . Now, imagine we are working in pixel space in the no-noise limit and we have a prior on the CMB that it is sky-uncorrelated, i.e. it is gaussian white noise. Then $\mathbf{C}(i, j) = \mathbf{c} \delta_{ij}$ say and $\hat{\mathbf{R}}$ reduces to $\mathbf{c}^{-1} \otimes \hat{\mathbf{x}}$, defining $\hat{\mathbf{x}} \equiv \sum_i \mathbf{X}(i) \mathbf{X}^T(i)$. Then $\hat{\mathbf{R}}^{-1}$ is just $\mathbf{c} \otimes \hat{\mathbf{x}}^{-1}$ using (B3). Eq. (28) simplifies to become independent of \mathbf{c} , and so, no matter what our prior on \mathbf{c} actually is, we have

$$\langle \mathbf{S}_{\text{CMB}}(i) \rangle = \mathbf{w}_{\text{s.p.}}^T \mathbf{X}(i) \quad (31)$$

with

$$\mathbf{w}_{\text{s.p.}} = \hat{\mathbf{x}}^{-1} \mathbf{e} (\mathbf{e}^T \hat{\mathbf{x}}^{-1} \mathbf{e})^{-1}. \quad (32)$$

For temperature alone, in which case \mathbf{e} is just $(1, \dots, 1)^T$, this is exactly the ILC result. So Eqs. (31) and (32) are the natural generalization of ILC when treating temperature and polarization in a unified manner. Typically the ILC procedure is described as choosing the linear combination of channels that has minimum variance whilst still retaining unit response to the CMB. Our Bayesian derivation here (albeit with flat priors on the foregrounds and a very strong, and obviously incorrect, prior on the CMB) here gives an alternative, more insightful, perspective. For example, we see that one might not need to correct for “cosmic covariance” [12, 35]; the ILC coefficients already give the mean maps. We might replace the “white noise CMB” prior with one based on a fiducial model up to an overall amplitude and thus derive “improved” ILC coefficients that correctly take into account spatial correlations in the CMB. A very simple case appropriate for low- l would be to take the CMB to be scale-invariant, $C_l \propto 1/(l(l+1))$, and work in harmonic space, calculating $\hat{\mathbf{R}}$ with the appropriate weights.

As mentioned below Eq. (30), we can calculate a measure of the foreground-induced pixel-pixel covariance of our maps. For clarity we shall do this for temperature alone in the no-noise limit. Then $\mathbf{S}_{\text{CMB}}(i)$ is simply $t(i)$, the temperature in the i 'th pixel. With the “white noise” CMB prior, \mathbf{c} is given by the single number ΔT^2 . With a Jeffreys' prior on ΔT^2 one finds

$$\langle \langle t(i)t(j) \rangle \rangle = \frac{1}{N_{\text{elem}} - n + 1} (\mathbf{X}(i) - \langle t(i) \rangle \mathbf{e})^T \frac{\hat{\mathbf{x}}^{-1}}{\mathbf{e}^T \hat{\mathbf{x}}^{-1} \mathbf{e}} (\mathbf{X}(j) - \langle t(j) \rangle \mathbf{e}). \quad (33)$$

This is a very reasonable result; uncertainties are given by a quadratic form on differences away from a blackbody, with the quadratic form determined by the covariance of the sample.

One can attempt generalizing the ILC procedure in other ways. For example, staying with temperature alone, one could keep the “white noise” CMB prior, but try to add in instrument noise at some level, taking it say to be sky-uncorrelated and isotropic, described by the same matrix \mathbf{n} at each pixel. Then the joint likelihood for \mathbf{w} and ΔT^2 , with a flat prior on ΔT^2 , turns out to be:

$$p(\Delta T^2, \mathbf{w} | \mathbf{X}) \propto \frac{\delta(\mathbf{w}^T \mathbf{e} - 1)}{\sqrt{\Delta T^2 + \mathbf{w}^T \mathbf{n} \mathbf{w}}} e^{-\frac{1}{2} \frac{\mathbf{w}^T \hat{\mathbf{x}} \mathbf{w}}{\Delta T^2 + \mathbf{w}^T \mathbf{n} \mathbf{w}}}, \quad (34)$$

and the joint maximum likelihood point for ΔT^2 and \mathbf{w} together occurs at the just same value for \mathbf{w} , namely $\hat{\mathbf{x}}^{-1} \mathbf{e} / (\mathbf{e}^T \hat{\mathbf{x}}^{-1} \mathbf{e})$, as in the noise-free case, with unchanged variance for \mathbf{w} . This is at odds with the idea of subtracting the noise contribution \mathbf{n} from the total variance $\hat{\mathbf{x}}$ (which might seem plausible on the grounds of focussing on the signal) in deriving the ILC coefficients.

VI. POSSIBLE EXTENSIONS, PRACTICAL CONSIDERATIONS AND CONCLUSION

While the scheme presented above treats foregrounds very generally, there is scope for further extension. One possibility is to allow for a slow variation of the mixing matrix across the sky. In fact it would be technically easiest to perform a low-order spherical harmonic expansion of the inverse pixel-space mixing matrix; then $\mathbf{w}^T(\mathbf{x})\mathbf{e} = \mathbf{i}$ would reduce to $\mathbf{w}_{00}^T \mathbf{e} = \mathbf{i}$ and $\mathbf{w}_{lm}^T \mathbf{e} = \mathbf{0}$ for $l > 0$. Ref. [15] applies an analogous approach to map making, allowing the ILC coefficients to vary across the sky. Notice that this is distinct from the harmonic approach of [13, 14] and is rather more physical when looked at from a component separation approach; the l -by- l analysis there effectively imagines the data from given direction to come from a convolution of signals around that direction, whereas [15] and the suggestion here consider the data in a given direction to come from signals in that direction alone.

Another natural variation would be to relax the flat prior on the foregrounds and rather try to marginalize over power spectra for the foregrounds. Now the separation achieved above between CMB and foregrounds would not occur and one would have to think about possible priors on the foreground part of the inverse mixing matrix (luckily we have seen that the base prior on \mathbf{a} rapidly becomes irrelevant as the number of elements increases). The underlying model would be close to that used in the Gibbs sampling approach of [29], and in the context of map making one would be left with a scheme very similar to that of [24] (but with the advantages of working with \mathbf{a}^{-1} rather than \mathbf{a}). Now one of the distinguishing features of the foreground signals is that they are non-gaussian, and one could hope to exploit this if one could develop a plausible non-gaussian correlated probability distribution for the foregrounds. As mentioned earlier, one might also consider non-gaussian corrections to the CMB itself, but in this case at least we suspect them to be small.

A conceptually simple change would be to relax the assumption that the mixing matrix \mathbf{a} is square, if say one is confident in the number of physical foreground emission processes operating. However, many of the manipulations of our approach depend technically on the invertibility of \mathbf{a} , so this change would be difficult to implement and would only be trustworthy if any neglected foreground components (e.g. coming from non-modelled spectral index variations) were below the level of the detector noise.

Let us now move on to discuss certain practical considerations in implementing a scheme such as that described here. An obvious difficulty with a straightforward implementation is the usual one confronting CMB analysis, namely the need to numerically invert large matrices corresponding to the large size of the data sets involved. Although the cosmic signal is naturally band limited, the cutoff is too high for a full resolution analysis over a good fraction of the sky to be practical.⁸ With simplifying assumptions however, some progress may be possible.⁹ In any case, a full resolution analysis is probably not even necessary. Instead, one might perform a split analysis, treating large and small angular scales in a different manner (see [36] for a discussion in the absence of foregrounds). While the large

⁸ Of course one might be able to apply the scheme directly to data from high resolution ground based experiments that only focus on small patches of the sky.

⁹ For example, one might work in harmonic space and assume that the detector noise, along with the signal, is diagonal in this basis. Then many of the matrix operations simplify. Still implementing a general cut using the projection technique, the computational burden might be reduced by the ratio of the cube of the number of modes projected out to the cube of the number of elements. A limited number of further corrections to the noise matrix might also be efficiently implementable using Woodbury formula techniques. Alternatively one might consider perturbative corrections to the noise model away from the (masked) isotropic case.

angular scales (where some of the most interesting new results might manifest themselves and where particularly B-mode polarization might have some signal relative to detector noise for Planck) might be dealt with accurately by an application of the method of this paper, an heuristic approximation to the method (or a different scheme altogether) might be all that is needed for high l . Of course the problem will not precisely factorize into a large scale part and a small scale part and so the coupling will have to be taken into account or shown to be negligible. In addition, there are issues with the production and interpretation of low resolution maps and their (inverse) noise matrices (see the description in [37] of the procedure used by the WMAP team), with difficulties due to aliasing and the finite size of the pixels themselves. In this work we have assumed that the signal has been deconvolved from the beam; the appropriate noise matrices must be obtained by transforming those appropriate to the common case of solving only for the beam-smoothed signal. Any beam uncertainties should be incorporated in the noise matrices also.

Often the time-ordered data is discretized into a map on the sphere using some pixelization scheme as opposed to going directly to harmonic space. Even if all signals are strictly band-limited, with non-zero spherical harmonic coefficients only for $l \leq l_{\max}$ say, fully describing such signals in pixel space requires more than $(l_{\max} + 1)^2$ pixels. One then has to be careful about the validity of certain matrix operations that need to be performed. For example, the inverse transformation from an unconstrained map in pixel space to harmonic space is not well-defined. Similarly, the pixel-space signal covariance matrix is non-invertible. The harmonic-space inverse noise matrix can be obtained by transforming the pixel-space inverse noise matrix though. An important issue arises if a galaxy cut projection is to be made: the mask itself needs to be band-limited in order for pixel-space and harmonic-space approaches to be equivalent.

While all of these complications need to be investigated and accounted for if necessary, it is not likely that they will lead to a fundamental problem in the scheme defined here for cosmic information extraction in the presence of foregrounds, just as they have not prevented standard cosmological analysis of CMB data.

In conclusion, this paper has described a prescription for the analysis of multiple sky maps for cosmological information in the presence of both detector noise and foregrounds. The prescription takes into account uncertainties in the foregrounds by modelling the foregrounds in a very general way and then marginalizing over the foregrounds. The noise can be very general, potentially correlated over the sky and even correlated between frequency channels, and include projections. Work is underway both to subject the prescription to extensive Monte Carlo testing and to apply it to the WMAP data. Assuming its performance is good the scheme should be ideal for the analysis of the temperature and polarization data from the forthcoming Planck satellite.

Acknowledgments

I thank George Efstathiou for directing my attention to the problem of the foreground contamination of the CMB, and thank him, Mark Ashdown, Anthony Challinor, Antony Lewis and Francesco Paci for helpful comments and discussions. I am supported by STFC.

APPENDIX A: CONSTRAINED SADDLE POINT INTEGRALS

This appendix explains the technique we use to evaluate constrained integrals in a saddle-point approximation.

Imagine we are performing an n -dimensional integral over variables x^i of some function $e^{-S(x)}$, with $m < n$ linearly-independent linear constraints of the form $c_i^k x^i = d^k$ (summed over i) on the x 's applied, i.e.

$$I = \int d^n x \delta(c_i^1 x^i - d^1) \dots \delta(c_i^m x^i - d^m) e^{-S(x)}. \quad (\text{A1})$$

First, we approximate the delta functions by a narrow gaussian, described with some covariance matrix D which will eventually be taken to zero:

$$\delta(c^T x - d) \equiv \delta(c_i^1 x^i - d^1) \dots \delta(c_i^m x^i - d^m) \sim \frac{1}{\sqrt{|2\pi D|}} e^{-(c_i^m x^i - d^m) D^{-1}{}_{mn} (c_i^n x^i - d^n)/2}. \quad (\text{A2})$$

Next, we find the saddle point of the entire integrand and approximate the integrand as a gaussian around the saddle point. The saddle point is where the first derivative

$$c D^{-1} (c^T x - d) \Big|_i + S_{,i} \quad (\text{A3})$$

of minus the exponent is zero, and the matrix M_{total} of second derivatives of minus the exponent has components

$$c D^{-1} c^T \Big|_{ij} + S_{,ij}. \quad (\text{A4})$$

We need to evaluate both exponent and prefactor terms to obtain the full saddle-point approximation to (A1). At the saddle point, it turns out that the “delta-function” part of the exponent tends to zero as $D \rightarrow 0$, as can be seen as follows. Multiplying (A3) by c on the left at the saddle point (“s.p.” in formulae below) and rearranging yields

$$c^T x - d = -D(c^T c)^{-1} S_{,i}|_{\text{s.p.}} \quad (\text{A5})$$

and so the “delta-function” part is

$$S_{,i}^T|_{\text{s.p.}} (c^T c)^{-1} D (c^T c)^{-1} S_{,i}|_{\text{s.p.}} / 2. \quad (\text{A6})$$

If there exists a sensible solution as $D \rightarrow 0$, then $S_{,i}|_{\text{s.p.}}$ will only have a weak dependence on D and so in the limit the above term will vanish.

The prefactor comes from the gaussian integral over deviations away from the saddle point, which is proportional to the reciprocal square root of M_{total} evaluated at the saddle point. Splitting M_{total} into the piece $cD^{-1}c^T$ from the delta function and the piece M with components $S_{,ij}|_{\text{s.p.}}$, we have $|M_{\text{total}}| = |M||D^{-1}||D + c^T M^{-1}c|$. Combining with the $\sqrt{|D|}$ term in the denominator of (A2) and taking the limit, we obtain an overall prefactor proportional to $1/\sqrt{|M||cM^{-1}c^T|}$.

Hence we obtain:

$$I \approx \frac{1}{\sqrt{|M/(2\pi)| |2\pi c^T M^{-1}c|}} e^{-S}, \quad (\text{A7})$$

with all quantities evaluated at the saddle point. Note that the result depends only on the original integrand and its second derivatives but evaluated at the saddle point of the combined integrand. Additionally,

$$\int d^n x \delta(c^T x - d) x e^{-S(x)} \approx x_{\text{s.p.}} I, \quad (\text{A8})$$

$$\int d^n x \delta(c^T x - d) x x^T e^{-S(x)} \approx \left(M^{-1} - M^{-1} c (c^T M^{-1} c)^{-1} c^T M^{-1} + x_{\text{s.p.}} x_{\text{s.p.}}^T \right) I. \quad (\text{A9})$$

We still have to actually find the saddle point and typically this can be done numerically with a Newton-Raphson-type approach as follows. With the first and second derivatives from formulae (A3) and (A4) above, the Newton-Raphson update is of the form:

$$\delta x = -M_{\text{total}}^{-1} (cD^{-1}(c^T x - d) + \nabla S(x)). \quad (\text{A10})$$

With the saddle point of the limit being the limit of the saddle point we can simply start from a point on the constraint surface and take the limit $D \rightarrow 0$ straight away to obtain the projected update formula:

$$\delta_{\text{proj}} x = - \left(M^{-1} - M^{-1} c (c^T M^{-1} c)^{-1} c^T M^{-1} \right) \nabla S(x) \quad (\text{A11})$$

(“proj” for projected). Iterating this will lead us to the saddle point, which can then be used in (A7) to obtain the approximate value for the integral.

In the special case that S is quadratic in x , the saddle point can be solved for analytically and is at $x = M^{-1} c (D + c^T M^{-1} c)^{-1} d$. So taking $D \rightarrow 0$, we have

$$x_{\text{s.p.}} = M^{-1} c (c^T M^{-1} c)^{-1} d \quad (\text{for } S \text{ quadratic in } x) \quad (\text{A12})$$

and hence

$$S_{\text{s.p.}} = d^T (c^T M^{-1} c)^{-1} d / 2 \quad (\text{for } S \text{ quadratic in } x). \quad (\text{A13})$$

Furthermore, Eqs. (A7), (A8) and (A9) become exact.

APPENDIX B: VECTORIZATION AND KRONECKER PRODUCTS

We here summarize vectorization of matrices and the Kronecker product of two matrices. See Appendix A of [38] for further identities and a recent discussion of vectorization in the context of CMB analysis.

The vectorization $\text{vec}(A)$ of an m -by- n matrix A is the mn -by-1 column vector formed by stacking the columns of A on top of each other. Explicitly, $\text{vec}(A)^T = (A_{11}, \dots, A_{m1}, A_{12}, \dots, A_{m2}, \dots, A_{1n}, \dots, A_{mn})$.

The Kronecker product of an m -by- n matrix A and a p -by- q matrix B is an mp -by- nq matrix, denoted $A \otimes B$, formed by appropriately stacking copies of the B matrix that have been multiplied by the elements of A :

$$A \otimes B = \begin{pmatrix} A_{11}B & A_{12}B & \cdots & A_{1n}B \\ A_{21}B & A_{22}B & \cdots & A_{2n}B \\ \vdots & \vdots & & \vdots \\ A_{m1}B & A_{m2} & \cdots & A_{mb}B \end{pmatrix}. \quad (\text{B1})$$

When the matrices are of compatible sizes such that the relevant products exist,

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD). \quad (\text{B2})$$

Hence,

$$(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}. \quad (\text{B3})$$

Also,

$$(A \otimes B)^T = A^T \otimes B^T. \quad (\text{B4})$$

A useful identity involving both vectorization and the Kronecker product is

$$\text{tr}(A^T BCD) = \text{vec}(A)^T (D^T \otimes B) \text{vec}(C). \quad (\text{B5})$$

APPENDIX C: SUMMARY OF MATRIX IDENTITIES USED

This paper makes extensive use of Sherman-Morrison/Woodbury-type formulae (see e.g. [39]) for matrix inverses and their determinants. Here we briefly show how these results may be derived in order to understand how and when they may be applied. We start by considering two related decompositions of the same blocked matrix:

$$\begin{pmatrix} a & -u \\ v^T & c^{-1} \end{pmatrix} = \begin{pmatrix} 1 & -uc \\ 0 & 1 \end{pmatrix} \begin{pmatrix} a + uc v^T & 0 \\ 0 & c^{-1} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ cv^T & 1 \end{pmatrix} \quad (\text{C1})$$

$$= \begin{pmatrix} 1 & 0 \\ v^T a^{-1} & 1 \end{pmatrix} \begin{pmatrix} a & 0 \\ 0 & c^{-1} + v^T a^{-1} u \end{pmatrix} \begin{pmatrix} 1 & -a^{-1}u \\ 0 & 1 \end{pmatrix}. \quad (\text{C2})$$

Taking the determinant of both decompositions and rearranging yields

$$|a + uc v^T| = |a| |c| |c^{-1} + v^T a^{-1} u|. \quad (\text{C3})$$

Inverting both decompositions and equating top-left corners yields

$$(a + uc v^T)^{-1} = a^{-1} - a^{-1} u (c^{-1} + v^T a^{-1} u)^{-1} v^T a^{-1}. \quad (\text{C4})$$

-
- [1] A. A. Penzias and R. W. Wilson, *Astrophys. J.* **142**, 419 (1965).
 - [2] R. H. Dicke, P. J. E. Peebles, P. G. Roll, and D. T. Wilkinson, *Astrophys. J.* **142**, 414 (1965).
 - [3] G. Hinshaw et al. (WMAP) (2008), 0803.0732.
 - [4] B. Gold et al. (WMAP) (2008), 0803.0715.
 - [5] M. R.olta et al. (WMAP) (2008), 0803.0593.
 - [6] J. Dunkley et al. (WMAP) (2008), 0803.0586.
 - [7] E. L. Wright et al. (WMAP) (2008), 0803.0577.
 - [8] R. S. Hill et al. (WMAP) (2008), 0803.0570.
 - [9] E. Komatsu et al. (WMAP) (2008), 0803.0547.

- [10] C. L. Bennett, G. F. Smoot, G. Hinshaw, E. L. Wright, A. Kogut, G. de Amici, S. S. Meyer, R. Weiss, D. T. Wilkinson, S. Gulkis, et al., *Astrophys. J. Lett.* **396**, L7 (1992).
- [11] C. Bennett et al. (WMAP), *Astrophys. J. Suppl.* **148**, 97 (2003), astro-ph/0302208.
- [12] G. Hinshaw et al. (WMAP), *Astrophys. J. Suppl.* **170**, 288 (2007), astro-ph/0603451.
- [13] M. Tegmark and G. Efstathiou, *Mon. Not. Roy. Astron. Soc.* **281**, 1297 (1996), arXiv:astro-ph/9507009.
- [14] M. Tegmark, A. de Oliveira-Costa, and A. Hamilton, *Phys. Rev.* **D68**, 123523 (2003), astro-ph/0302496.
- [15] J. Kim, P. Naselsky, and P. R. Christensen (2008), 0803.1394.
- [16] A. Hyvärinen, J. Karhunen and E. Oja, *Independent Component Analysis* (John Wiley and Sons, 2001).
- [17] D. Maino, S. Donzelli, A. J. Banday, F. Stivoli, and C. Baccigalupi (2006), astro-ph/0609228.
- [18] D. Maino et al. (2001), astro-ph/0108362.
- [19] D. Maino, A. J. Banday, C. Baccigalupi, F. Perrotta, and K. M. Gorski, *Mon. Not. Roy. Astron. Soc.* **344**, 544 (2003), astro-ph/0303657.
- [20] S. Donzelli et al., *Mon. Not. Roy. Astron. Soc.* **369**, 441 (2006), astro-ph/0507267.
- [21] A. Bonaldi et al. (2007), 0707.0469.
- [22] A. Bonaldi, L. Bedini, E. Salerno, C. Baccigalupi, and G. De Zotti, *Mon. Not. Roy. Astron. Soc.* **373**, 263 (2006), astro-ph/0609701.
- [23] R. Vio and P. Andreani (2008), 0802.0400.
- [24] J. Delabrouille, J. F. Cardoso, and G. Patanchon, *Mon. Not. Roy. Astron. Soc.* **346**, 1089 (2003), astro-ph/0211504.
- [25] G. Patanchon, J. F. Cardoso, J. Delabrouille, and P. Vielva, *Mon. Not. Roy. Astron. Soc.* **364**, 1185 (2005), astro-ph/0410280.
- [26] J.-F. Cardoso, M. Martin, J. Delabrouille, M. Betoule, and G. Patanchon (2008), 0803.1814.
- [27] H. K. Eriksen et al., *Astrophys. J.* **641**, 665 (2006), astro-ph/0508268.
- [28] R. Stompor, S. Leach, F. Stivoli, and C. Baccigalupi (2008), 0804.2645.
- [29] H. K. Eriksen et al. (2007), 0709.1058.
- [30] H. K. Eriksen et al. (2007), 0709.1037.
- [31] H. K. Eriksen et al., *New Astron. Rev.* **50**, 861 (2006), astro-ph/0604160.
- [32] The Planck Collaboration (2006), astro-ph/0604069.
- [33] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).
- [34] G. Efstathiou, *Mon. Not. Roy. Astron. Soc.* **346**, L26 (2003), astro-ph/0306431.
- [35] L.-Y. Chiang, P. D. Naselsky, and P. Coles (2007), 0711.1860.
- [36] G. Efstathiou, *Mon. Not. Roy. Astron. Soc.* **349**, 603 (2004), astro-ph/0307515.
- [37] N. Jarosik et al. (WMAP), *Astrophys. J. Suppl.* **170**, 263 (2007), astro-ph/0603452.
- [38] S. Hamimeche and A. Lewis (2008), 0801.0554.
- [39] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes: The Art of Scientific Computing* (Cambridge University Press, 1986).